

Distributed Odor Source Localization

Adam T. Hayes, Alcherio Martinoli, Rodney M. Goodman

Abstract—This paper presents an investigation of odor localization by groups of autonomous mobile robots. First, we describe a distributed algorithm by which groups of agents can solve the full odor localization task. Next, we establish that conducting polymer-based odor sensors possess the combination of speed and sensitivity necessary to enable real world odor plume tracing, and we demonstrate that simple local position, odor, and flow information tightly coupled with robot behavior is sufficient to allow a robot to localize the source of an odor plume. Finally, we show that elementary communication among a group of agents can increase the efficiency of the odor localization system performance.

Keywords—Chemical Plume Tracing, Olfaction, Odor Localization, Distributed Sensing, Collective Autonomous Robotics, Swarm Intelligence

I. INTRODUCTION

RECENT advances have been made in understanding biological odor localization and tracking as developed in moths [1], [2] and rats [3] in the air, and lobsters [4] and stomatopods [5] in water. Biology utilizes olfaction for a wide variety of tasks including finding others of the same species, communication, behavior modification, avoiding predators, and searching for food. Animals use a combination of ‘hardware’ (frequency of receptor adaptation, perhaps), ‘software’ (temporal integration and/or spatial integration), and behavioral search strategies (both intrinsic and landmark-based) to locate odor sources. Odor localization is in essence a behavioral problem that varies from animal to animal. While some animals exploit fluid information at different layers (lobster) or several residues on the ground (ants), others can track odors in the air (moths) or use a combination of information (dogs). From an engineering standpoint there are advantages to combining odor tracking with mobile robots [6], such as in the detection of chemical leaks and the chemical mapping of hazardous waste sites. We are interested in developing groups of small mobile robots that use odor tracking algorithms, multiple sensory modalities (e.g. odometry, anemometry, olfaction), and sensory fusion to search out and identify sources of odor.

We design the agent control algorithms using principles of Swarm Intelligence (SI), a computational and behavioral metaphor for solving distributed problems that takes its inspiration from biological examples provided by social insects. In most biological cases studied so far, robust and capable group behavior has been found to be mediated by nothing more than a small set of simple interactions among individuals and between individuals and the environment [7]. The application of SI principles to autonomous collec-

tive robotics aims to develop robust task solving by minimizing the complexity of the individual units and emphasizing parallelism, exploitation of direct or indirect interactions, and distributedness. These principles favor the design of behavior-based robotic systems, similar at the individual level to those of Brooks [8] and Arkin [9], which emphasize tight coupling between sensation and action, avoidance of representational knowledge, and action decomposition into contextually meaningful units [9]. The main advantages of the SI approach are three: first, scalability from a few to thousands of units, second, flexibility, as units can be dynamically added or removed without explicit reorganization, and third, increased system robustness, not only through unit redundancy but also through the design of minimalist units. Several examples of collective robotics tasks solved with SI principles can be found in the literature: aggregation [10] and segregation [11], beacon localization [12], stick pulling [13], and collective transportation [14].

The aim of the case study described in this paper is four-fold. First, we describe a distributed algorithm by which groups of agents can solve the full odor localization task. Second, we establish that conducting polymer-based odor sensors possess the combination of speed and sensitivity necessary to enable real world odor plume tracing. Third, we demonstrate via real robots and embodied simulations that simple sensory information tightly coupled with robot behavior is sufficient to allow an agent to find the source of an odor plume. Last, we show that integrating the information collected by a group of agents in an elementary manner can increase the efficiency of the odor localization system performance.

II. THE ODOR LOCALIZATION PROBLEM

A. Task Description

The general odor localization problem addressed in this paper is as follows: find a single odor source in an enclosed 2D area as efficiently as possible. This can be broken down into three subtasks: plume finding - coming into contact with the odor, plume traversal - following the odor plume to its source, and source declaration - determining from odor acquisition characteristics that the source is in the immediate vicinity. Plume finding amounts to a basic search task, with the added complication, due to the stochastic nature of the plume, that a simple sequential search is not guaranteed to succeed. Plume traversing requires more specialized behavior, both to progress in the direction of the source and to maintain consistent contact with the plume. Source declaration does not necessarily have to be done using odor information, as typically odor sources can be sensed via other modalities from short range, but here we

propose a solution using no extra sensory apparatus.

As an odor source dissolves into a fluid medium, an odor plume is formed. The turbulent nature of fluid flow typically breaks the plume into isolated packets, areas of relative high concentration surrounded by fluid that contains no odor [15]. The task of odor localization is thus one of plume traversal, or following the trail of odor packets upstream to the source. This becomes difficult as odor packets become more sparse (due to source intermittency and diffusion below detectable levels) and more dispersed (due to flow meander).

Previous odor localization research has utilized concentration gradient information to locate the centerline of a plume, and then either anemotaxis [16], [17] or further chemotaxis [18] to proceed to the source. These systems were restricted to operation in the proximal region of the plume (within 2 m of the source) and had to move slowly (.01-.03 m/s) so that concentration gradient information could be extracted with reasonable accuracy using sensors with sub-Hz response or recovery times. Although these efforts were successful in demonstrating the feasibility of odor localization with mobile robots, it is not clear that any method that involves spatial concentration extraction will extend to more sparse plumes (i.e. longer plume tracking distances), since as odor information becomes less frequent, concentration integration times will increase, decreasing system performance accordingly [19]. Moreover, these investigations were limited by the speed of the sensors they incorporated. In the distal plume region where plume information is intermittent, sensors that are too slow to register passing odor packets are of little use.

B. Biological Inspiration

Although the approach of moving slowly and continually sampling odor and flow data to reduce environmental noise is used in nature (starfish) and has been applied to robotic systems [17], [20], environmental and behavioral constraints (e.g. significant plume sparseness or meander, time critical performance) can render these systems ineffective. In that case, upon sensing an odor signal, a good policy is to move directly upwind, as a good immediate local indication of source direction under such circumstances is the instantaneous direction of flow [21]. When the odor is no longer present, a good strategy is to perform a local search (known as casting in the biological literature) until it is reacquired, as the location of the previous packet encounter provides the best immediate estimate of where the next will occur. This type of surge-cast behavior has been observed in moths [22], and its performance has been studied in simulation [2].

The previous work on this odor localization algorithm was aimed at studying biology, which limited the sensory and behavioral time scales investigated. When applying these ideas to artificial systems, however, the separation between algorithm and underlying hardware is much more clear, and it no longer makes sense to constrain behavior strictly by sensory response characteristics. Therefore, in this work key aspects of the search behavior, such as surge

duration and casting locality, are treated as algorithm parameters.

C. The Spiral Surge Algorithm

The basic odor localization algorithm used in this study, Spiral Surge (SS), is shown in Figure 1. It consists of different behaviors related to the three different subtasks.

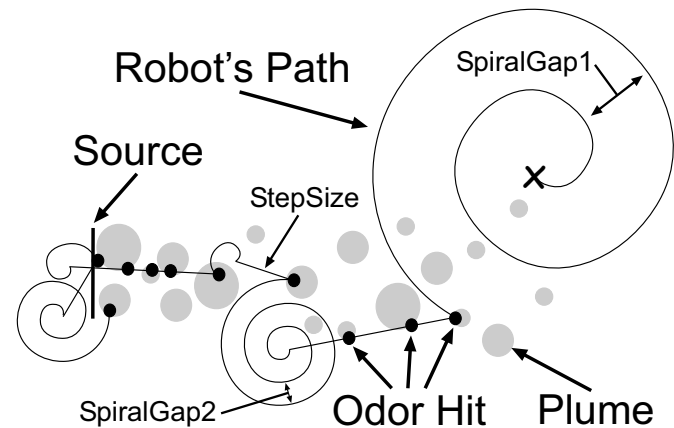


Fig. 1. Spiral Surge odor localization behavior.

TABLE I
SPIRAL SURGE ALGORITHM PARAMETERS

PARAMETER	DESCRIPTION
SPIRALGAP1	Initial spiral gap width
SPIRALGAP2	Plume reacquisition spiral gap width
STEP SIZE	Surge distance post odor hit
CAST TIME	Length of time before reverting from reacquisition to initial search spiral
SRCDEC THRESH	Significance threshold between consecutive separate odor hits
SRCDEC COUNT	Number of significant differences before source declaration

Plume finding is performed by an initial outward spiral search pattern (SPIRALGAP1). This allows for thorough coverage of the local space if the total search area is very large and initial information can be provided by the deployment point (an external 'best guess' as to source location). Alternatively, if no a priori knowledge is available, a spiral with a gap much greater than the arena size (producing essentially straight line search paths) provides an effective, although not optimal [23], search procedure.

Plume traversal is performed using a type of surge algorithm. When an odor packet is encountered during spiraling, the robot samples the wind direction and moves upwind for a set distance (STEP SIZE). If during the surge another odor packet is encountered, the robot resets the surge distance but does not resample the wind direction. After the surge distance has been reached, the robot begins a spiral casting behavior, looking for another plume hit.

The casting spiral can be tighter than the plume finding spiral (SPIRALGAP2), as post surge the robot has information about packet density and a thorough local search is a good strategy. If the robot subsequently re-encounters the plume, it will repeat the surging behavior, but if there is no additional plume information for a set amount of time (CASTTIME), the robot will declare the plume lost and return to the plume finding behavior (with a wider, less local, spiral gap parameter).

Source declaration can be accomplished using the fact that a robot performing the plume traversal behavior at the head of a plume will tend to surge into an area where there is no plume information, and then spiral back to the origin of the surge before receiving another odor hit. If the robot keeps track internally of the post spiral inter-hit distances (using odometry, for example, which is sufficient because information must be accurate only locally), a series of small differences can indicate that the robot has ceased progress up the plume, and must therefore be at the source. However, because small inter-hit distances can occur in all parts of the plume, this method is not foolproof, and tuning of the difference threshold (SRCDECTHRESH), as well as the number of observed occurrences before source declaration (SRCDECCOUNT), is required to obtain a particular performance within a given plume. See Table I for a summary of individual SS parameters.

SS uses only binary odor information generated from a single plume sensor because this is the most simple and reliable type of information that can be obtained from real hardware in the temporal operating regime of interest. There may be information encoded in distal fine plume structure [24], however, due to the highly stochastic nature of turbulent fluid flow and the odor-packet nature of the plume, it is unclear that more complex sensing – via graded intensity information or larger fixed sensor arrays – would benefit an odor localizing agent when flow information is available through other means.

D. Collaborative Spiral Surge

While more complex odor sensing may be beneficial to the odor localization task, another possible route to greater efficiency is physical distribution of the odor sensing elements, which in principle could improve system speed and robustness via parallelization of the search procedure. This can be achieved by constructing an arbitrarily large and complex single robot or, perhaps more conveniently, distributing a number of sensors throughout a group of smaller, more simple communicating robots. With a suitable command and control interface, this collective can be viewed as an ‘odor localization sensor’ in much the same way a single larger robot, or more generally device, could. One way to increase the performance of such a robot swarm is collaboration between individual nodes. In particular, if collaboration is obtained with simple explicit communication schemes such as binary signaling, the team performance can be enhanced without losing autonomy or significantly increasing complexity at the individual level.

Several simple types of communication can be integrated

into basic SS. In this study, we examine the performance impact of three types of communication: no communication (NONE), a ‘come here’ signal emitted by upwind surging robots that causes all robots downwind or with no plume information to surge in the direction of the calling robot (ATTRACT), and a ‘stop’ signal emitted by the first robot to receive odor information that causes all other robots to surge away from the signaling robot and then enter a power save mode from which they cannot be awakened (KILL). We investigate the influence of these types of communication across group size to determine their impact on system efficiency.

E. Odor Localization Performance

To study odor localization, we place groups of agents within a starting area inside an enclosed arena containing an odor plume. Over repeated trials we measure the time and distance traveled by the whole group until the group completes the task, which can be defined in a number of ways: an agent comes within a given radius of the plume source (allocentric determination – useful to emulate a non-odor related target sensor that each robot might carry), an agent declares the plume source found (egocentric determination, no additional sensor necessary), or any combination and extension (i.e. multiple declarations required within a given radius) thereof. In this paper, we examine both single robot allo- and egocentric stopping criteria, so for the purposes of performance we simply assume that some measure of time and group energy (which can be considered proportional to the sum of the individual distances traveled) necessary for task completion exist (T_{TC} , D_{TC}).

Efficiency for the odor localization task cannot be defined in the general case. Instead, we combine the two basic measures of task performance in a task specific manner. Since these measures are physically independent, a composite metric incorporating a particular weighting of these two basic factors can be considered.

$$Q = \alpha T_{TC} + \beta D_{TC} \quad (1)$$

$$P = \frac{\alpha T_{MIN} + \beta D_{MIN}}{Q} \quad (2)$$

Q is an arbitrary weighting of time and distance. By choosing specific values for α and β , the appropriate relationship can be generated for evaluating any particular application. The form of P ensures that for any α and β greater than 0, the optimal system will achieve a performance of 1, and any that require more time or distance will have a performance less than 1. We determine the optimum values for the given task (T_{MIN} , D_{MIN}) from an agent executing the optimal behavior (a straight line path from start to goal areas at maximum speed). Maximum speed, which determines the relationship between the time and distance values, is determined by the maximum safe operating speed of the agent in the given environment. In this study we set α and β so that the time and energy components of the task factor equally into the minimum cost, so $\frac{\alpha}{\beta} = \frac{D_{MIN}}{T_{MIN}}$.

III. MATERIALS AND METHODS

A. Real Robots

We use Moorebots, as shown in Fig. 2, which were originally designed by Owen Holland at the University of West England, Bristol, U.K.. Each 24 cm diameter robot is equipped with two DC motor-driven wheels, a castor wheel, a 2 Mbit wireless LAN transceiver, and 12-bit A/D and D/A converters. See [25] for a more detailed robot description. We extended this basic configuration with four infra-red range sensors for collision avoidance, a single odor sensor, and a hot wire anemometer. On-board high-level control is provided by a PC104 based Intel 386 processor running Linux. Low level control such as motor speed regulation is executed by dedicated hardware interfaced to the PC104 bus.




Fig. 2. A Moorebot equipped with wind, odor, and proximity sensors, as well as markings for overhead tracking.

B. Robot Arena and Infrastructure

The plume traversal arena is 6.7 by 6.7 m. The odor plume is created by a 23 cm square hot water pan and a bank of 5 fans 30 cm in diameter, and it extends diagonally from one corner of the arena toward the opposite corner. Flow characteristics based on data taken along the plume axis 15 cm above the floor are summarized in Table II. The coefficient of variation is a measure of the intensity of the flow turbulence. It represents the ratio of the standard deviation of the wind velocity to the mean wind velocity, and 20% is a value typically measured outdoors [16].

TABLE II
WIND FIELD CHARACTERIZATION

Distance from source [m]	1	4	8
Mean wind speed [m/s]	1.13	1.01	.34
Coefficient of variation [%]	15.4	21.2	52.0

The robot start area is located in the corner opposite the plume source. An overhead camera tracking system, combined with a radio LAN among the robots and an external

workstation, is used to log position data during the trials, reposition the robots between trials, and emulate the binary communication signals. Trials of different group size are interleaved and inactive robots are automatically positioned at recharging stations. The arena layout, as seen from the overhead camera, is shown in Figure 3.




Fig. 3. Real robot arena as seen from overhead camera.

C. Odor Sensor

While many types of odor sensing technology currently exist [26], a good combination of ease of transduction, reversibility, reproducibility, tunability, ease of production, robustness across environments, miniaturization, and speed is offered by carbon-doped polymer sensors [27]. This odor sensor detects the presence of an airborne substance through a change in the electrical resistance of a chemically sensitive carbon-doped polymer film [28]. While this type of sensor can lack baseline stability, it is very fast (response times $< .1$ s [20]), and signal processing techniques can be used to counteract its baseline drift.

We fabricate sensors from solutions consisting of 20% carbon black and 80% polymer (poly-vinylpyrrolidone) dissolved in dichloromethane, using methods as described in [29]. The conducting polymer solution is spray coated [30] onto the surface of a surface mount universal board so that the sensor film closes the circuit between two mounting pads. Polymer solution is applied until sensor resistance nears 100 kOhm, and baseline resistances typically settle to a value between 30 and 300 kOhm after a 24 hour drying period. A sensor close-up can be seen in Figure 4.

The interface circuitry applies an input bias voltage across a multiplexer selectable range resistor to generate a current through the sensor via a Wilson current source. The output voltage across the sensor is then filtered to remove high frequency noise and buffered for reading. The variable bias voltage and selectable range resistor allow a wide range of sensor baseline resistances (10 kOhm to 10 MOhm) and automatic calibration, an important feature because polymer sensors are difficult to fabricate precisely and their baselines drift over time. The calibration procedure consists of switching through all range resistors with




Fig. 4. Odor sensor close-up.

the bias voltage centered (and no stimulus present), choosing the resistor that results in an output closest to the desired output, and then adjusting the bias voltage until the desired baseline output is achieved. The resistor and bias values are then stored for later use. The desired output value is 25% of the ADC's range, as the sensor values are more likely to drift up than down.

Previous versions of the interface circuit used a local analog feedback loop to maintain the output voltage at a constant level. However, this low-pass hardware filtering attenuated not only the sensor drift but the signal as well, reducing sensor sensitivity. Sensitivity is crucial for the study of plume tracing, because the agents must be able to sense a meaningful plume structure, not simply respond when very close to the odor source. In our indoor experimental set-up, room ventilation is limited, so enhancement of the plume signal is not an option. Thus, instead of using analog feedback, we digitally filter the output signal and record an odor hit whenever the filtered signal rises above some threshold. We use a sixth order Butterworth bandpass filter, and we set the filter parameters by comparing the power spectral density given no stimulus with the output power spectral density with the robot stationary in the distal part of the plume. Given that we want to select a frequency range that provides the highest possible SNR, we use bandpass cutoff values of .3 to 1.8 Hz based on the data shown in Figure 5. Although the sensors can respond at higher frequencies, no information is available above 2 Hz given our transduction circuitry and experimental conditions. The amplitude threshold for odor detection is set at 4 times the baseline standard deviation (recorded from 10,000 samples taken at an average rate of 85 Hz following calibration) to render false positives improbable.

When executing the odor localization algorithm, the odor sensor polling rate averages 85 Hz. Because the robot CPU is performing the polling, the filtering, and handling all other tasks the robot requires (e.g. communications, high-level motor control, and memory management), the sensor polling rate is not precise, and we do not use a real




Fig. 5. Power spectral density of the odor sensor output when no stimulus is present and when the robot is in the distal end of plume.

time Linux kernel (which could provide reliable, although slower, polling rates) due to the overhead it requires. We do not account for this imprecision in the digital filter, and treatment of the polling jitter, through, for example, the use of a dedicated microcontroller to take sensor readings, could increase sensitivity. However, the combination of the calibration procedure and digital filtering produces a robust binary odor detection sensor. Figure 6 compares raw and filtered data from the distal end of the plume against filtered baseline data from the same sensor. The detection threshold is plotted 4 std above 0, and the raw data has been DC shifted about -3 V for ease of presentation. The presence of odor hits 8 m from the plume source shows that a significant plume stimulus exists to be tracked, even in the distal plume region where odor information is intermittent. Mapping the plume using a random walk behavior indicates that the plume is stable over time and across robots (see Figure 7 and Figure 8).

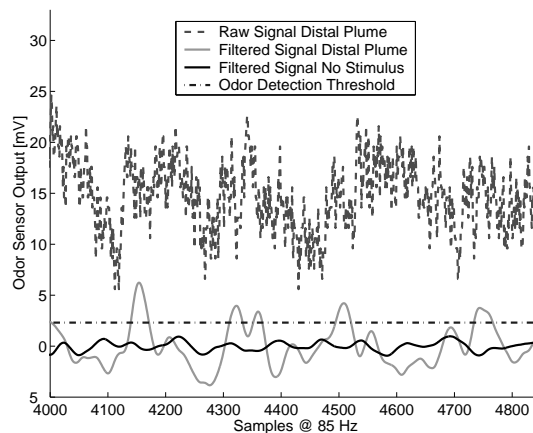


Fig. 6. Raw distal plume data, filtered distal plume data, and filtered baseline data. The threshold is 4 std above 0.




Fig. 7. Total plume hits received by 6 real robots over 1 hour while performing a random walk behavior.




Fig. 8. Plume hits received by 6 individual real robots over 1 hour while performing a random walk behavior.

D. Wind Sensor

The anemometer is a Shibaura F6201-1 air flow sensor, as used by [16], which can sense wind flow down to .05 m/s. It is enclosed in a tube to provide unidirectional sensitivity, which, combined with a scanning behavior, allows the robot to measure wind direction. When wind direction information is required, the robot first rotates 90 degrees, then rotates slowly 360 degrees while reading the wind sensor output, and finally rotates back to the heading corresponding to the highest sensor value. The robot takes the shortest path back to the desired heading, and either over or under rotates to the target to account for the 1 s time delay of the internal anemometer processing circuitry. The initial rotation reduces the probability that the robot begins facing upwind, in which case the discontinuity in the scanning behavior can degrade the resulting wind direction value. Wind sensor performance has yet to be fully characterized due to the requirements of a suitable testing environment (flow must be laminar), although the data from the odor localization experiments suggests it is sufficient for the given task. A wind map of 2102 individual

samples averaged spatially is shown in Figure 9.




Fig. 9. Average wind direction in plume traversal arena as measured by the real robots. Plume source at upper right. Arrow lengths are proportional to the uniformity of flow direction at the tail of each arrow.

E. Embodied Simulation

When studying the performance of distributed robotic systems, it can be useful to model the system using different levels of abstraction. Successful modeling provides a way of understanding the essential aspects of the system, as well as a significantly decreased evaluation time, which enables a more complete investigation of the system parameter space. Models also allow treatment of environmental conditions which (for some technical limitation) cannot be implemented physically. In this work, the use of a model permits us to enlarge the search arena and examine agent performance on the full odor localization problem. Because the source declaration phase of the task can lead to elevated agent densities around the source, and thus is very sensitive to inter-agent repulsion parameters, non-embodied simulations, which can only approximate such interactions, are not able to provide faithful results. Thus, we use Webots [31], a 3D sensor-based, kinematic simulator, originally developed for Khepera robots [32], to systematically investigate the performance of SS in simulation. This embodied simulator has previously been shown to generate data that closely matches real Khepera [13], [10], [33] and Moorebot [12] experiments, so we are confident that real robot behavior is accurately captured.

We performed simulations in an arena modeled after the physical arena, as shown in Figure 10, to verify that our simulator produces accurate results, and we also used a 25 times (area) larger arena, which is large enough to allow study of the full odor localization problem (see Figure 11). The agent behavioral algorithms correspond exactly to those used by the real robots. To properly capture the plume stimulus, we incorporated a series of leaky source 2D PLIF plume images generated in a water flume by Donald Webster and Philip Roberts at Georgia Tech [34] [35]. Such ‘plume movies’, even though they do not capture the influ-




Fig. 10. Webots plume traversal arena with average plume intensity map.




Fig. 11. Layout of larger Webot arena.

ence of the agents on plume dynamics, offer a good approximation to the discretized (packet-like) nature of odor stimulus received in real environments. We scaled the recorded plume data to imitate the average speed and envelope of the real plume data (see Figure 12 and Figure 7), and tuned the odor sensitivity threshold (higher threshold leads to less odor information) based on performance observed in our real arena. Odor hit frequency differences between the real and simulated maps are due to the fact that for efficiency the simulated sensors are bandwidth limited only by the update rate of the plume data (10 Hz) rather than by a bandpass filter like the one used on the real robots (.3-1.8 Hz). While we used the same plume stimulus in both arenas, simulations of the physical arena incorporate flow information taken directly from the real robot data (as shown in Figure 9), and in the larger arena, in an effort to emulate more open flow patterns, wind information is generated by adding $\pm 10\%$ white noise to a constant direction parallel to the main plume axis.




Fig. 12. Plume hits received by 6 simulated robots over 1 hour.

IV. RESULTS AND DISCUSSION

A. Real Robots

The real robots experiments focus on the plume traversal subtask because it contains most of the plume related complexity present in the full odor localization task, and due to experimental limitations it is not feasible to study all phases with real robots at this time. Since source declaration is not being studied, a trial is complete when a robot reaches a given distance, the source find radius, from the plume source. To justify the high density of agents in the plume (which would be unlikely given that in the general problem the plume area is a small percentage of the total search area), we allow ATTRACT communication between the agents to hold the group together as it traverses the plume.

We tested real robot plume traversal performance using two sets of SS parameters and two control experiments. Only SPIRALGAP2 and STEPSIZE are considered because we are looking only at the plume traversal aspect of the task. SS1 represents a non-local search in that its search paths are straight and its surges extend to the boundaries of the arena. SS2 uses a smaller spiral gap and surge length to perform a more local exploration of the arena. **Random Odor** uses SS2 parameters, and receives odor hits that are generated from the time sequence of SS2 odor hits but are not correlated with robot position in the arena. This control experiment investigates whether an algorithm incorporating precise odor packet location information is more efficient than a blind upwind surging behavior. An alternative experiment could be to decouple the wind source from the odor source by creating a wind field with an array of fans, but due to practical limitations in our experimental set-up, the **Random Odor** case was easier to implement and provided equivalent information from a proof-of concept point of view. **Random Walk** takes straight line paths and random avoidance turns at boundaries (using no odor or flow information) to provide a traversal performance baseline. Specific parameters relating to the real robot tests

are listed in Table III. 15 trials of each group size were run for SS1, SS2 and Random Odor, and 30 trials were run for Random Walk due to the high variance of performance values. All error bars in the plots represent standard error.

TABLE III
PLUME TRAVERSAL PARAMETER VALUES

Agent speed	.325 m/s
Arena length	6.7 m
Plume length	8 m
Plume speed	~ 1 m/s
Source find radius	.88 m
Plume: Arena area	1:2.3
Goal: Search perimeter	1:18.0
T_{MIN}	19.0 s
D_{MIN}	6.2 m
$\frac{\alpha}{\beta}$.326 [m/s]
SS1: SPIRALGAP2	1785 km
SS1: STEPSIZE	9.1 m
SS2: SPIRALGAP2	.357 m
SS2: STEPSIZE	.91 m

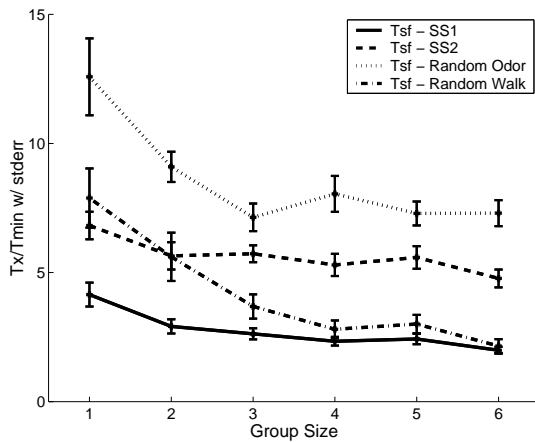


Fig. 13. Normalized time across group size for real robot trials. Lower values are better.

Figures 13 and 14 show that for all conditions studied, traversal time decreases with group size while group distance traveled increases. This indicates, as expected for a search task, that as time becomes more important to performance than energy usage, larger group sizes will be preferred.

Figure 15 shows that while single robots are generally most efficient in this arena (given this particular choice of α and β), SS1 generates the best results for each group size (significant via K-S test to $p < .01$ for group size $\in \{1, 2, 3\}$), demonstrating successful real robot plume traversal. Random Odor performs worse than SS2 for all group sizes (significant as above for group size $\in \{1, 2, 4, 6\}$), indicating that location of odor information is an important aspect of the search algorithm. This means




Fig. 14. Normalized distance across group size for real robot trials. Lower values are better.




Fig. 15. Performance across group size for real robot trials. Higher values indicate better performance.

that SS is actually plume tracing rather than simply localizing the source of the wind, because if it were only wind localizing, one would expect Random Odor to perform exactly the same as SS2. Also, SS2 performs worse than SS1 (significant as above for all group sizes), suggesting that local search is not a good strategy in this small arena where the goal-to-search perimeter ratio is high (i.e., it is likely to find the goal by chance). The Random Walk behavior retains relatively constant performance across group size, and at the larger group sizes its performance tends to approach the optimal observed performance. This suggests that as a search arena becomes overcrowded, random movement becomes the best strategy.

B. Embodied Simulations

We successfully reproduced the real robot performance data in Webots, as shown in Figure 16. Data represents 1000 trials per group size. All parameters in Table III apply to the Webots data as well. Only SS1 for group size of one robot produces significantly different results (as de-

terminated by a 2-tailed K-S test with $p < .01$) between Webots and the real robots, and even in this case the error bars overlap. Because our Webots data closely matches our available real robot data, it is reasonable that further simulated experiments will accurately reflect real world behavior.




Fig. 16. Performance of real robot and Webots trials across group size. Higher values indicate better performance.

C. The Full Odor Localization Task

The principal limitation of the experiments described thus far is the relatively small arena available for the real robots. In simulation we can expand the arena size and move the start area outside the plume extent. This enables the study of all phases of the odor localization task and calls for a change to the task stopping condition. Source declaration defines the end of a trial, and the time and distance data below contains only trials that result in a successful source declaration, i.e., a declaration of the source within the source found radius.

Techniques are under development to optimize system performance across the entire array of SS parameters, but for the purposes of this work, to illustrate that a distributed group of sensors can confront the odor localization problem and show that simple communication can affect performance, a functional set will suffice. The SS parameters are based on SS2, because in comparison to SS1 its tight casting spirals are more likely to result in small inter-hit distances. We assign SPIRALGAP1 to a large value to generate straight line search paths, set CASTTIME to a behaviorally reasonable value, and fix the source declare parameters in a functional regime. For the sake of simplicity, all communication signals are assumed to extend throughout the testing arena. Environmental and algorithmic parameter values that differ from the real robots experiments are shown in Table IV.

We examine the performance impact of the three types of communication described earlier: NONE, ATTRACT, and KILL, which correspond to parameter sets None, Attract, and Kill. 1500 trials were performed for each parameter

TABLE IV
FULL TASK PARAMETER VALUES (SIMULATION)

Arena Length	33.6 m
Plume: Arena Area	1:58
Goal: Search Perimeter	1:55
Wind Noise	$\pm 10\%$
SPIRALGAP1	3570 m
SPIRALGAP2	.357 m
STEP SIZE	.91 m
CAST TIME	96 s
SRCDEC THRESH	.27 m
SRCDEC COUNT	3

set and group size.




Fig. 17. Normalized time and distance to find source across group size and communication type.

Figure 17 shows the time and distance necessary for each group of robots to find the source of the plume. Note that qualitatively the curves resemble those of Figures 13 and 14, as one would expect, except here the reduction in time with group size is more pronounced. This is a result of the increased significance of the search phase of the task, which benefits more directly from the parallel nature of distributed random search than the other phases. In fact, the traverse and declare phases are possibly done best serially, as Kill, which uses only one agent for these tasks, requires shorter group distances than the other sets. However, note that Kill needs more time than None to complete the task as group size grows, which is a reflection on the high temporal cost of plume loss when the parallel plume-reacquisition search capability is lost. Attract does not seem to offer any benefits at all, as it uses the same amount of time as Kill but much more energy. This is due to a high interference rate as the agents collect in the plume, impeding in particular the proper spiral search paths necessary for source declaration.

Figure 18 indicates that Kill is the most efficient form of communication for this environment, with its substan-




Fig. 18. Performance across group size and communication type. Higher values are better.

tial energy savings outweighing its slightly longer running times. Also, again for this environment, the efficiency peaks at a group size of 5 agents. In fact, all of the communication types peak at group sizes larger than 1, which suggests that distributing and coordinating sensors in this manner is an effective way to increase system efficiency. It is important to note that these results, particularly the relative efficiencies of the different communication types, are likely to be heavily dependent on the task description. More complex plume stimuli (containing a sparser signal or more large scale meander) that are more difficult for individual robots to track should favor systems that achieve higher agent densities in the plume. This is difficult to test currently because real plume data with large scale meander are not available, although this is an avenue of ongoing research.




Fig. 19. Median distance from source upon declaration across group size and communication type.

Figure 19 shows the effect of group size on declaration accuracy across plume type. All communication types are able to yield high success rates (>95%) for low group sizes,

with only None's performance dropping off to below 70% at the large group sizes. Kill maintains a constant performance across group size, as should be expected because only one robot is performing the operation. More interestingly, None degrades at larger group sizes, while Attract does not. This can be explained if the agents in Attract, even though their physical interactions are repulsive, are able to maintain a tightly clustered group at the head of the plume through their communication signals. Thus even though physical interactions slow down the declaration process, when source declaration does occur the declaring agent is likely to be near the plume source. However, because None also has repulsive physical interactions but no attractive force other than the plume itself, the cloud of agents formed around the head of the plume will be less dense, and when there are many agents the probability of the declaring agent being far from the source increases. It is important to note that these results depend heavily on the details of the physical interactions between the agents, which are accurately modeled by the embodied simulator, although interactions between the plume and the agents (which are not modeled) may also play a role.




Fig. 20. Normalized time and distance across group size for different subtasks for Kill.

Figure 20 offers a detailed look at Kill performance during each task phase. Time to find the plume is defined as the length of time from the start of the trial to the first odor plume encounter. Time to traverse the plume is defined as the amount of time between the first odor encounter and a robot entering within the source find radius of the source. Time to declare the plume is the time between entering the source find region and declaring that the plume has been found. Note that for any particular trial it need not be the same robot completing each phase. The time to find the plume decreases with group size, yet the group distance remains nearly constant, indicating that the random search is effectively parallel. Also note that at high numbers of agents, because they all start in the same location, a performance penalty appears due to unnecessary search overlap near the start area. Since only a single robot is performing

the traversal and declaration tasks for all group sizes, it is not surprising that the time required does not change, and the linear increase in group distance can be attributed to the non-zero energy consumption of the inactive robots.

V. CONCLUSION

This paper presented an investigation of odor localization by groups of autonomous mobile robots. First, we described a distributed algorithm by which groups of agents can solve the odor localization task. Because this algorithm is based upon both odor and flow information, it is not designed to function in environments in which flow is too weak to detect reliably (typically $< .05m/s$ [19]). Still, there are a broad range of military and industrial situations that involve stronger flows (in particular any outdoor environment) for which it does apply.

Next, because we were able to show that our robots could detect plume information as far as 8 m away from the source, we established that conducting polymer-based odor sensors possess the combination of speed and sensitivity necessary to enable real world odor plume tracing. This is important because previous efforts at mobile robotic odor localization [17] relied on slower sensor technologies which in turn restricted the classes of algorithms that could be applied. Yet through an appropriate combination of calibration and filtering techniques we showed it is possible to overcome the lack of stability inherent in polymer odor sensors and investigate new algorithm domains permitted by fast response times and low-power, lightweight (i.e. potentially mobile) odor sensors. Note we are not proposing that polymer sensors are the only suitable sensor type for all environments and robotic applications. Rather, these findings suggest that response times should be factored in along with steady-state sensitivity when selecting a sensor for a particular task, and sensor precision may have diminished importance because even binary odor concentration information can be useful for an odor localization system.

We also demonstrated that simple sensory information tightly coupled with robot behavior is sufficient to allow a robot to find the source of an odor plume. This shows the power of integrating actuation into sensory systems, and suggests that complicated sensory transduction may not be necessary when a behaving sensory mechanism is well tuned to its designated task [36]. In addition, we showed that integrating the information collected by a group of agents in an elementary manner can increase the efficiency of the odor localization system performance, an avenue that has not been previously explored using real robots. If we view the entire system as an odor localization sensor, the distributed approach opens up a new axis of optimization (inter-agent communication) not available when only a single unit is considered, and the organizational principles of SI allow such distributed systems to remain scalable and require minimal additional complexity. The particular communication types explored in this paper represent the most basic interactions available, and as the complexity of the task description increases (more complicated plume types, higher frequencies of false-positive odor hits), corre-

spondingly more complicated interaction schemes (greater number of signals, variable signaling range) will likely be necessary to yield a performance benefit. More work needs to be done to determine how this complex parameter space can be explored in a systematic way.




Fig. 21. An Alice robot with 400 element olfaction chip.

Finally, it may seem contradictory that while the SI approach stressed in this paper emphasizes minimalism, the actual robots used in this study feature general purpose microprocessors and high bandwidth communication. However, because care was taken to keep system requirements low, the algorithms used in this study can be ported directly to much less expensive or smaller platforms, such as the Alice robot [37] shown in Figure 21. Only when robot swarms can be implemented on a large scale will the robust nature of these systems be fully exploited. As more advanced sensors become available which can combine sensitivity, discrimination, and mobility, such as the polymer odor sensing matrix consisting of 400 elements integrated on a chip [38] shown next to the Alice, truly useful real-world odor localization systems will become feasible.


ACKNOWLEDGMENTS

We would like to thank Owen Holland, Sanza Kazadi, Jim Pugh, Robert Enright, Ladd Van Tol, and Andrew Lundsten for their contributions on various hardware and software aspects of the project. We would also like to thank our collaborators in the DARPA-ONR Chemical Plume Tracing Program for their valuable input. This work is supported in part by the Center for Neuromorphic Systems Engineering as part of the National Science Foundation Engineering Research Center Program under grant EEC-9402726. Support has also been received from the DARPA Dog's Nose program under grant DAAK60-97-K-9503. This work is also supported in part by the Office of Naval Research under grant N00014-98-1-0821 'Chemical Plume Tracing Via Biological Inspiration and Electronic Olfaction', and by the Army Research Office under MURI grant DAAG55-98-1-0266 'Understanding Olfaction: From Detection to Classification'. Adam Hayes is supported by


a National Science Foundation Graduate Research Fellowship.

REFERENCES


- [1] R. T. Carde and A. Mafra-Neto, "Effect of pheromone plume structure on moth orientation to pheromone," in *Perspectives on Insect Pheromones. New Frontiers*, R. T. Carde and A. K. Minks, Eds., pp. 275–290. Chapman and Hall, N.Y., 1996.
- [2] J. H. Belanger and M. A. Willis, "Adaptive control of odor guided locomotion: Behavioral flexibility as an antidote to environmental unpredictability," *Adaptive Behavior*, vol. 4, pp. 217–253, 1996.
- [3] U. Bhalla and J. M. Bower, "Multi-day recording from olfactory bulb neurons in awake freely moving rats: Spatial and temporally organized variability in odorant response properties," *J. of Computational Neuroscience*, vol. 4, pp. 221–256, 1997.
- [4] J. Atema, "Eddy chemotaxis and odor landscapes: Exploration of nature with animal sensors," *Biological Bull.*, vol. 191, pp. 129–138, 1996.
- [5] M. J. Weissburg, "From odor trails to vortex streets: Chemo and mechanosensory orientation in turbulent and laminar flows," in *Orientation and Communication in Arthropods*, M. Lehrer, Ed., pp. 215–246. Birkhauser, Basel, 1997.
- [6] R. A. Russell, *Odor detection by mobile robots*, World Scientific, Singapore, 1999.
- [7] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, New York, US, 1999.
- [8] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Robotics and Automation*, vol. RA-2, pp. 14–23, March 1986.
- [9] R. C. Arkin, *Behavior Based Robotics*, MIT Press, Cambridge, MA, 1998.
- [10] A. Martinoli, A. J. Ijspeert, and F. Mondada, "Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots," *Robotic and Autonomous Systems*, vol. 29, pp. 51–63, 1999.
- [11] O.E. Holland and C. Melhuish, "Stigmergy, self-organization, and sorting in collective robotics," *Artificial Life*, vol. 5, pp. 173–202, 1999.
- [12] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Comparing distributed exploration strategies with simulated and real autonomous robots," in *Proc. of the fifth Int. Symp. on Distributed Autonomous Robotic Systems DARS-2000*, L. E. Parker, G. Bekey, and J. Barhen, Eds., Knoxville, Tennessee, October 2000, pp. 261–270, Springer Verlag.
- [13] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella, "Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment," *Autonomous Robots*, vol. 11, no. 2, pp. 149–171, 2001.
- [14] C. R. Kube and E. Bonabeau, "Cooperative transport by ants and robots," *Robotics and Autonomous Systems*, vol. 30, pp. 85–101, 2000.
- [15] J. Murlis, J. S. Elkington, and R. T. Carde, "Odor plumes and how insects use them," *Annu. Rev. Entomol.*, vol. 37, pp. 505–532, 1992.
- [16] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi, "Odor-source localization in the clean room by an autonomous mobile sensing system," *Sensors and Actuators B*, vol. 33, pp. 115–121, 1996.
- [17] T. Nakamoto, H. Ishida, and T. Moriizumi, "A sensing system for odor plumes," *Analytical Chemistry*, vol. 71, no. 15, pp. 531A–537A, August 1999.
- [18] R. A. Russell, D. Thiel, R. Deveza, and A. Mackay-Sim, "A robotic system to locate hazardous chemical leaks," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Nagoya, 1995, pp. 556–561.
- [19] H. Ishida, T. Nakamoto, T. Moriizumi, T. Kikas, and J. Janata, "Plume-tracking robots: A new application of chemical sensors," *Biological Bulletin*, vol. 200, pp. 222–226, April 2001.
- [20] S. Kazadi, R. Goodman, D. Tsikata, and H. Lin, "An autonomous water vapor plume tracking robot using passive resistive polymer sensors," *Autonomous Robots*, vol. 9, no. 2, pp. 175–188, 2000.
- [21] C. T. David, J. S. Kennedy, J. S. Ludlow, and J. N. Perry, "A re-appraisal of insect flight towards a point source of wind-borne odor," *Journal of Chemical Ecology*, vol. 8, pp. 1207–1215, 1982.
- [22] N. J. Vickers and T. C. Baker, "Reiterative responses to single strands of odor promote sustained upwind flight and odor source location by moths," *Proceedings of the National Academy of Sciences USA*, vol. 91, pp. 5756–5760, 1994.
- [23] D. W. Gage, "Randomized search strategies with imperfect sensors," in *Proceedings of SPIE Mobile Robots VIII*, Boston, September 1993, vol. 2058, pp. 270–279.
- [24] D. R. Webster, S. Rahman, and L. P. Dasi, "On the usefulness of bilateral comparison to tracking turbulent chemical odor plumes," *Limnology and Oceanography*, vol. 46, no. 5, pp. 1048–1053, 2001.
- [25] A.F.T. Winfield and O.E. Holland, "The application of wireless local area network technology to the control of mobile robots," *Microprocessors and Microsystems*, vol. 23, pp. 597–607, 2000.
- [26] H. T. Nagle, R. Guitierrez-Osuna, and S. S. Schiffman, "The how and why of electronic noses," *IEEE Spectrum*, vol. 35, no. 9, pp. 22–31, September 1998.
- [27] M. C. Lonergan, E. J. Severin, B. J. Doleman, S. A. Beaver, R. H. Grubbs, and N. S. Lewis, "Array-based vapor sensing using chemically sensitive, carbon black-polymer resistors," *Chem. Mater.*, vol. 8, pp. 2298–2312, 1996.
- [28] M. S. Freund and N. S. Lewis, "A chemically diverse conducting polymer-based electronic nose," *Proceedings of the National Academy of Sciences USA*, vol. 92, pp. 2652, 1995.
- [29] B. J. Doleman, M. C. Lonergan, E. J. Severin, Vaid T. P., and Lewis N. S., "Quantitative study of the resolving power of arrays of carbon black-polymer composites in various vapor-sensing tasks," *Anal. Chem.*, vol. 70, pp. 4177–4190, 1998.
- [30] A. J. Matzger, C. E. Lawrence, R. H. Grubbs, and N. S. Lewis, "Combinatorial approaches to the synthesis of vapor detector arrays for use in an electronic nose," *J. Comb. Chem.*, vol. 2, pp. 301–304, 2000.
- [31] O. Michel, "Webots: Symbiosis between virtual and real mobile robots," in *Proceedings of the First International Conference on Virtual Worlds, VW'98*, Paris, France, July 1998, pp. 254–263, Springer Verlag.
- [32] F. Mondada, E. Franzi, and P. lenne, "Mobile robot miniaturization: A tool for investigation in control algorithms," in *Proc. of the Third International Symposium on Experimental Robotics ISER-93*, T. Yoshikawa and F. Miyazaki, Eds., Kyoto, Japan, 1993, pp. 501–513, Springer Verlag.
- [33] A. Martinoli, A. J. Ijspeert, and L. G. Gambardella, "A probabilistic model for understanding and comparing collective aggregation mechanisms," in *Proc. of the Fifth Int. European Conf. on Artificial Life ECAL-99, Lecture Notes in Computer Science*, D. Floreano, F. Mondada, and J.-D. Nicoud, Eds., pp. 575–584. Springer Verlag, Lausanne, Switzerland, September 1999.
- [34] D. R. Webster and M. J. Weissburg, "Chemosensory guidance cues in a turbulent chemical odor plume," *Limnology and Oceanography*, vol. 46, no. 5, pp. 1034–1047, 2001.
- [35] D. R. Webster, S. Rahman, and L.P. Dasi, "Laser-induced fluorescence measurements of a turbulent plume," *ASCE Journal of Engineering Mechanics*, To appear (2002).
- [36] B. Webb, "View from the boundary," *Biological Bulletin*, vol. 200, pp. 184–189, April 2001.
- [37] G. Caprari, P. Balmer, R. Piguat, and R. Siegwart, "The autonomous microrobot 'alice': a platform for scientific and commercial applications," in *Proc. of the Ninth Int. Symp. on Micromechanics and Human Science*, Nagoya, Japan, November 1998, pp. 231–235.
- [38] J. A. Dickson and R. M. Goodman, "Integrated chemical sensors based on carbon black and polymer films using a standard CMOS process and post-processing," in *Proc. of the ISCAS2000, 2000 IEEE International Symposium on Circuits and Systems, Emerging technologies for the 21st century*, Geneva, Switzerland, May 2000, pp. 341–344.



Adam T. Hayes received an A.B. Magna Cum Laude in Computer Science from Harvard University in 1998, and a Ph.D. in Computation and Neural Systems from the California Institute of Technology in 2002. His dissertation focused on a methodology for designing self-organized robotic systems that emphasizes task implementation and employs machine learning techniques to assist system development. This design process has been applied to the task of odor localization to demonstrate its capabilities. His current research interests include the development of biologically motivated sensory systems that enable autonomous robots to function in dynamic environments.



Alcherio Martinoli received a M. Sc. in Electrical Engineering from the Swiss Federal Institute of Technology in Zurich (ETHZ) in 1992, and a Ph.D. in Computer Science from the Swiss Federal Institute of Technology in Lausanne (EPFL) in 1999. From 1992 to 1994, he spent one year as Research Scientist at the Institute of Biomedical Engineering of the ETHZ and one year at the Institute of Industrial Automation of the Spanish Research Council in Madrid, Spain. His Ph.D. focused on modeling and performance prediction of distributed, mobile, robotic systems, evolutionary methods for designing and optimizing distributed control algorithms, and development of collective-specific tools for experiments with miniature robots. He is currently Senior Research Fellow and head of the Collective Robotics Group at the California Institute of Technology, Pasadena, U.S.A. His research interests include swarm intelligence, collective robotics, distributed control for embedded systems, and machine-learning optimization techniques applied to distributed problems.



Rodney M. Goodman received his B.Sc. degree with Honors in Electrical Engineering from Leeds University, Yorkshire, UK in 1968, and his Ph.D. degree in Electronics at the University of Kent at Canterbury, UK in 1976. From 1975 to 1985 he was a member of the faculty of the Department of Electrical and Electronic Engineering at the University of Hull, UK. In 1985 Dr. Goodman joined the faculty of the California Institute of Technology where he was Professor of Electrical Engineering until

September 2001. Dr Goodman left Caltech in October 2001 to focus on his entrepreneurial activities, and is currently a consultant to several advanced technology start-up companies in the Pasadena area, including Cyrano Sciences, the electronic nose company, of which he is a founder. In addition, Dr. Goodman is currently Faculty Associate in Electrical Engineering at Caltech. Dr. Goodman's current research interests are in intelligent information processing systems, electronic nose technology, and distributed communications networks of sensors and actuators. In particular, novel control architectures for multiple autonomous mobile robots, and machine consciousness are being pursued.